

Michael Rescorla

Book review: Gualtiero Piccinini // physical computation

**Article (Published version)
(Non-refereed)**

Original citation:

Rescorla, Michael (2016) *Book review: Gualtiero Piccinini // physical computation*. British Journal for the Philosophy of Science - Review of Books . ISSN 0007-0882

© 2016 The Author

This version available at: <http://eprints.lse.ac.uk/73691/>

Available in LSE Research Online: April 2017

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

Gualtiero Piccinini // Physical Computation

 bjpsbooks.wordpress.com/2016/09/07/gualtiero-piccinini-physical-computation/

Reviewed by Michael Rescorla

Physical Computation: A Mechanistic Account

Gualtiero Piccinini

Oxford: Oxford University Press, 2015, £35 (hardback)

ISBN 9780199658855

Gualtiero Piccinini's *Physical Computation: A Mechanistic Account* develops a systematic theory of computing systems. On display throughout are virtues familiar from Piccinini's previous writings on the topic: an admirably straightforward writing style; a synoptic perspective on computation's varied aspects, as studied within logic, computer science, artificial intelligence, cognitive science, engineering, neuroscience, and physics; extensive knowledge of computing practice; and a gift for conveying that knowledge in accessible, non-technical terms.

As the title indicates, Piccinini advances a *mechanistic* theory of physical computation. On his approach, a physical computing system is a physical mechanism that has the function of processing vehicles in accord with rules. We can categorize computations through properties of the vehicles and the rules. For example, digital computation manipulates digits, which are unambiguously distinguishable by the computing mechanism under normal circumstances. Analogue computation operates over continuous variables, which are not unambiguously distinguishable because they can only be measured within a certain margin of error. Piccinini develops his mechanistic account in truly impressive detail, offering systematic elucidations of 'mechanism', 'function', 'rule', 'vehicle', and other key concepts. He critiques rival accounts, and he applies his own account to sundry topics of great interest (for example, the Church–Turing thesis, pancomputationalism, the computational theory of mind). The result is an imposing framework, presented with tremendous verve and skill, that sheds light upon numerous facets of computation.

Piccinini's approach will remind some readers of Marcin Milkowski's *Explaining the Computational Mind* ([2013]), which also analyses computation in mechanistic terms. However, the two books differ considerably. For example, Milkowski embraces the thesis that computation is information-processing, whereas Piccinini regards that thesis warily. It would have been useful for Piccinini to compare his view more systematically with Milkowski's. Such a comparison might have helped readers better assess how the two views relate to one another.

One of Piccinini's main targets is a view that he calls the 'semantic account of computation'. According to the semantic account, every computing system has semantic or representational properties. Jerry Fodor advocates the semantic account, which he summarizes with the memorable slogan 'no computation without representation' ([1975], p. 34). Against the semantic account, Piccinini protests that many computing systems do not have anything like semantic or representational properties as standardly construed by philosophers. For example, a Turing machine might manipulate meaningless strings of digits. We can fully understand this machine's computations without citing anything like a semantics for the strings.

I agree with Piccinini that the semantic account of computation is implausible. Still, I worry that his treatment downplays crucial representational aspects of numerous computations. Representation plays a vital role within computer science, cognitive science, and most other disciplines that study computation. For example:

- Computer scientists often describe a computer as executing arithmetic operations, such as addition or division. A computer can add or divide numbers only if it can represent numbers. So these descriptions presuppose that the computer's states have representational properties. Only by citing representational properties of the computer's states can we explain why we built the computer in the first place: to execute appropriate arithmetical operations.
- Cognitive scientists often individuate mental states in representational terms. For example, perceptual

psychology describes how the perceptual system transits from proximal sensory stimulations to perceptual states that estimate shapes, sizes, colours, and other properties of distal objects (Burge [2010]; Rescorla [2015]). Any description along these lines individuates perceptual states through representational relations to specific distal shapes, sizes, colours, and other such properties. Representational description of perception has proved explanatorily fruitful, yielding rich understanding of computations executed by the perceptual system.

In these examples, and many others, one can fully understand the relevant computation only when one considers its representational properties. Piccinini acknowledges as much (p. 49). In practice, though, he devotes little space to representational aspects of computation. He focuses primarily on non-representational, mechanistic aspects. Readers who take Piccinini as their guide will emerge with scant appreciation for the pervasive role that representation plays within the study of artificial and natural computing systems.

Piccinini argues that even when representational description illuminates a computation, mechanistic description enjoys causal or explanatory priority. This seems plausible for many (perhaps all) artificial computing systems. Something like Piccinini's favoured mechanistic descriptions play a central role within computer science practice. Such descriptions illuminate the design, construction, and manipulation of artificial computing systems. However, matters are less straightforward when we consider computational modelling of the mind. In this context, it is not evident that Piccinini-style mechanistic description plays a valuable role.

The crucial notion here is medium-independence: 'concrete computations and their vehicles can be defined independently of the physical media that implement them' (p. 122). A computation may in principle be implemented in any arbitrary physical medium, such as silicon chips, neurons, wood levers-and-pulleys, and so on. What matters is the system's functional organization, not its physical constitution. Thus, medium-independence is a strong form of multiple realizability in Putnam's ([1967]) sense. As Piccinini notes, most processes studied within science are not medium-independent. For example, human digestion is not medium-independent, because it hinges on specific biochemical properties of enzymes, organs, and so on. However, Piccinini holds that mechanistic computation is medium-independent.

The medium-independent viewpoint seems plausible when applied to artificial computing systems. Computer scientists frequently prescind from medium-specific details of physical machines. They frequently offer abstract descriptions compatible with innumerable possible hardware implementations. Ultimately, of course, we must adduce medium-specific physical details. Otherwise, we will not be able to build the machine or repair it when it breaks. For many purposes, though, medium-specific physical details do not matter. Computer scientists do not usually care about a physical computer's hardware properties, so long as those properties enable the machine to execute desired computational operations (for example, adding two numbers together or moving a symbol from one memory register to another).

Piccinini's medium-independent viewpoint seems less plausible when applied to the mind. Piccinini thinks that cognitive science should strive to isolate multiply realizable, non-representational descriptions of mental computation. Yet such descriptions play little or no role in numerous areas of cognitive science. For example, current cognitive science assigns no significant role to multiply realizable, non-representational descriptions of perception (Rescorla [2015]) or motor control (Rescorla [2016]). Nor do I see there any reason why future cognitive science should follow Piccinini's advice and isolate medium-independent mechanistic descriptions of these phenomena. I agree with Piccinini that we should investigate how representational mental activity is mechanically implemented. But there is no evident reason why a satisfying mechanistic account should be medium-independent (Rescorla [forthcoming]).

The holy grail for cognitive science is to illuminate how mental activity is implemented in one very specific medium: the brain. How do neurons encode beliefs, desires, intentions, and other mental states? How does neural tissue implement characteristic mental processes such as perception and motor control? Answering these questions requires careful attention to medium-specific details about neurons, synapses, and the like. I see no reason why a good theory must also include a medium-independent mechanistic component. As applied to many core mental processes, I suspect that Piccinini's favoured mechanistic level is an idle wheel that adds no explanatory value.

Tellingly, Piccinini provides no concrete examples where his mechanistic approach tallies with successful scientific explanation of mental activity. Throughout the book, he draws his examples almost exclusively from the study of artificial computing systems. In lieu of detailed case studies involving mental computation, Piccinini offers a highly abstract argument that representational description of computation always assumes non-representational description in the background. As he puts it, 'semantic individuation presupposes non-semantic individuation' (p. 33). However, Piccinini's argument does not have the right form to show that mental computations fall under explanatorily significant medium-independent mechanistic descriptions. At best, the argument shows that representational description of mental computation is grounded in non-representational description. The argument leaves open that representational description of mental computation is grounded in non-representational neurophysiological description, without any significant role for non-representational medium-independent description.

Even taken on its own terms, the argument is problematic. At a crucial juncture, the argument enjoins us to reject 'primitivism about content, according to which we can simply posit contents without attempting anything like a reductive analysis of what it is for states to have content' (p. 35). Piccinini launches two objections against primitivism. First, 'contents are not the kind of thing that can be posited without a reductive account; they are not basic enough; they are not similar enough to fermions and bosons' (p. 35). Second, primitivism 'flies in the face of scientific practice' (p. 36), especially computer science and cognitive science.

I find both objections unconvincing. Reductive analysis for core scientific concepts is relatively rare. For example, we have nothing like reductive analyses for virus, interest rate, and force. Unreduced concepts do not inhibit good explanation in biology, economics, or physics. Why should they inhibit good explanation in psychology? Specifically, why does the concept 'representation' need reductive analysis in order to figure in good psychological explanations? Despite what Piccinini suggests, many areas of cognitive science happily traffic in representational notions without any attempt at reductive analysis. For example, perceptual psychology freely cites representational properties of perceptual states, without even gesturing towards an account of what it is for a perceptual state to have some representational property (for example, what it is for a perceptual state to represent a distal shape).

Piccinini is hardly alone in demanding a reductive analysis of representation. Often, the demand is a residual legacy of behaviourism, especially as filtered through Quine's ([1960]) virulent hostility to representational notions. (Piccinini (p. 36) cites behaviourism as an antecedent to his anti-primitivist analysis.) There prevails in certain philosophical circles a conviction that representation is unscientific, unclear, or even spooky, its explanatory bona fides to be securely established only by reducing it to the non-representational. In my opinion, the conviction has not yet been supported by any compelling arguments. It also conflicts with our best cognitive science, which offers thoroughly representational treatments of numerous core mental phenomena.

Piccinini says that cognitive science should make itself 'mechanistically legitimate' by supplementing representational description with non-representational computational description (p. 36). He offers his account as a template for the requisite non-representational mechanistic descriptions. I have questioned how well the proposed template fits scientific research into many core mental phenomena. Nevertheless, Piccinini's discussion is a notable contribution that offers a bounty of insights into computation and computing practice. All philosophers interested in computation must read this highly informative and thought-provoking book.^[1]

Michael Rescorla
Department of Philosophy
University of California, Los Angeles
rescorla@ucla.edu

References

Burge, T. [2010]: *Origins of Objectivity*, Oxford: Oxford University Press.

Fodor, J. [1975]: *The Language of Thought*, New York: Thomas Y. Crowell.

Milkowski, M. [2013]: *Explaining the Computational Mind*, Cambridge, MA: MIT Press.

Putnam, H. [1967]: 'Psychophysical Predicates', in W. Capitan and D. Merrill (eds), *Art, Mind, and Religion*, Pittsburgh: University of Pittsburgh Press.

Quine, W. V. [1960]: *Word and Object*, Cambridge, MA: MIT Press.

Rescorla, M. [2015]: 'Bayesian Perceptual Psychology', in M. Matthen (ed.), *The Oxford Handbook of the Philosophy of Perception*, Oxford: Oxford University Press.

Rescorla, M. [2016]: 'Bayesian Sensorimotor Psychology', *Mind and Language*, **31**, pp. 3–36.

Rescorla, M. [forthcoming]: 'From Ockham to Turing—and Back Again', in A. Bokulich and J. Floyd (eds), *Turing 100: Philosophical Explorations of the Legacy of Alan Turing*, Dordrecht: Springer.

Notes

[1] Thanks to Gualtiero Piccinini for many helpful discussions of these issues over the years. Thanks also to Dimitri Coelho Mollo for comments that improved this review.